

### EXAMINER'S AMENDMENT

An examiner's amendment to the record appears below. Should the changes and/or additions be unacceptable to applicant, an amendment may be filed as provided by 37 CFR 1.312. To ensure consideration of such an amendment, it MUST be submitted no later than the payment of the issue fee.

Authorization for this examiner's amendment was given in a telephone interview with Donald Gibson on 6/9/09.

The application has been amended as follows:

1. (Currently Amended) A system for allocating a resource in a network comprising: multiple network data ports to connect to the network to receive and route requests to servers in the network; first logic coupled to the network data ports for determining if any of a plurality of persistence policies, comprising at least one first persistence policy and at least one second persistence policy, is applicable to a service request from a client in the network, and, if so, allocating the resource to the request based on application of the persistence policy determined to be applicable, wherein the first persistence policy, if applicable, directs the service request to a server, previously or currently connected to the client, that is identified based on content of the service request, the content comprising server, session or cookie information within the service request or corresponding packet at one or more layers corresponding to OSI layers 5-7, and wherein the second persistence policy, if applicable, directs the service request to a server, previously or currently connected to the client, that is identified based on client information within the service request or corresponding packet; the first logic configured so that: when a service request for a content-enabled service, comprising a request for a connection to a server at

a layer corresponding to any of OSI layers 5-7, is received, the first logic determines if the at least one first persistence policy is applicable; when a service request for a non-content-enabled service, comprising a request for a connection to a server at a layer corresponding to OSI layer 4, is received, the first logic determines if the at least one second persistence policy is applicable; when the service request for a content-enabled service is received but it is determined that the at least one first persistence policy is inapplicable, the first logic determines if the at least one second persistence policy is applicable; wherein the first logic is shared by and supports service requests for both content-enabled and non-content enabled services; and second logic coupled to the network data ports for allocating the resource to the request based on application of a load balancing policy only if none of the plurality of persistence policies is determined to be applicable as determined by the first logic; wherein, when the service request for a content-enabled service is bound to a server identified through application of a load balancing policy, such request is bound to the server only after multiple persistence policies are checked for applicability; and one or more of the foregoing elements are implemented, embodied, maintained or controlled by one or more processor devices.

3. (Currently Amended) A system for allocating a resource, in a network, to a resource request, the request having an originator, based on application of a persistence policy comprising: multiple network data ports to connect to the network to receive and route requests to servers in the network; first logic coupled to the data ports for determining whether an allocation exists or recently expired for the originator of the resource request, based on application of any of a plurality of persistence policies, comprising at least one first persistence

policy and at least one second persistence policy, to the request, and, if so, identifying the resource which is the subject of the existing or recently expired allocation, wherein the first persistence policy, if applicable, directs the resource request to a server, currently or previously connected to the originator of the request or corresponding client, that is identified based on content of the resource request, the content comprising server, session or cookie information within the resource request or corresponding packet at one or more layers corresponding to OSI layers 5-7, and wherein the second persistence policy, if applicable, directs the resource request to a server, currently or previously connected to the originator of the request or corresponding client, that is identified based on client information within the request or corresponding packet; the first logic configured so that: when a resource request for a content-enabled service, comprising a request for a connection to a server at a layer corresponding to any of OSI layers 5-7, is received, the first logic determines whether an allocation exists or recently expired for the originator of the resource request by applying the at least one first persistence policy, or the at least one second persistence policy when the at least one first persistence policy is determined to be inapplicable; when a resource request for a non-content-enabled service, comprising a request for a connection to a server at a layer corresponding to OSI layer 4, is received, the first logic determines whether an allocation exists or recently expired for the originator of the resource request by applying the at least one second persistence policy; wherein the first logic is shared by and supports resource requests for both content enabled and non-content enabled services; and second logic in the network for allocating the resource, once identified, to the resource request; wherein, when the service request for a content-enabled service is bound to a server identified through application of a load balancing policy, such request is bound to the server only

Art Unit: 2444

after multiple persistence policies are checked for applicability; and one or more of the foregoing elements are implemented, embodied, maintained or controlled by one or more processor devices.

20. (Currently Amended) A method of allocating a resource in a network comprising:

providing multiple data ports to connect to the network to receive and route requests to servers in the network; determining if any of a plurality of persistence policies, comprising at least one first persistence policy and at least one second persistence policy, is applicable to a service request from a client in the network, and, if so, allocating the resource to the request based on application of the persistence policy determined to be applicable, wherein the first persistence policy, if applicable, directs the service request to a server, currently or previously connected to the client, that is identified based on content of the service request, the content comprising server, session or cookie information within the service request or corresponding packet at one or more layers corresponding to OSI layers 5-7, and wherein the second persistence policy, if applicable, directs the service request to a server, currently or previously connected to the client, that is identified based on client information within the service request or corresponding packet; when a service request for a content-enabled service, comprising a request for a connection to a server at a layer corresponding to any of OSI layers 5-7, is received, the determining step comprises determining if the at least one first persistence policy is applicable; when a service request for a non-content-enabled service, comprising a request for a connection to a server at a layer corresponding to OSI layer 4, is received, the determining step comprises determining if the at least one second persistence policy is applicable; when the service request for a content-enabled service is received, but it is determined that the at least one first persistence

policy is inapplicable, the determining step comprises determining if the at least one second persistence policy is applicable; performing the foregoing steps through logic coupled to the data ports that is shared by and supports service requests for both content enabled and non-content enabled services; and allocating the resource to the request based on application of a load balancing policy only if none of the persistence policies in the plurality of persistence policies is determined to be applicable in the foregoing determining step; wherein, when the service request for a content-enabled service is bound to a server identified through application of a load balancing policy, such request is bound to the server only after multiple persistence policies are checked for applicability; and one or more of the foregoing steps are performed in, by or for one or more processor devices.

21. (Currently Amended) A method of allocating a resource, in a network, to a resource request based on application of a persistence policy, the request having an originator, comprising: providing multiple network data ports to connect to the network to receive and route requests to servers in the network; determining whether an allocation exists or recently expired for the originator of the resource request, based on application of any of a plurality of persistence policies, comprising at least one first persistence policy and at least one second persistence policy, to the request, and, if so, identifying the resource which is the subject of the existing or recently expired allocation, wherein the first persistence policy, if applicable, directs the resource request to a server, currently or previously connected to the originator of the resource request or corresponding client, that is identified based on content of the resource request, the content comprising server, session or cookie information within the resource request or corresponding

Art Unit: 2444

packet at one or more layers corresponding to OSI layers 5-7, and wherein the second persistence policy, if applicable, directs the resource request to a server, currently or previously connected to the originator of the request or corresponding client, that is identified based on information within the resource request or corresponding packet other than content; when a resource request for a content-enabled service, comprising a request for a connection to a server at a layer corresponding to any of OSI layers 5-7, is received, the determining step comprises determining whether an allocation exists or recently expired for the originator of the resource request by applying the at least one first persistence policy, or the least one second persistence policy when the at least one first persistence policy is determined to be inapplicable; when a resource request for a non-content-enabled service, comprising a request for a connection to a server at a layer corresponding to OSI layer 4, is received, the determining step comprises determining whether an allocation exists or recently expired for the originator of the request by applying the at least one second persistence policy; performing the foregoing steps through logic coupled to the data ports that is shared by and supports both resource requests for content enabled and non-content enabled services; and allocating the resource in the network, once identified, to the resource request; wherein, when the service request for a content-enabled service is bound to a server identified through application of a load balancing policy, such request is bound to the server only after multiple persistence policies are checked for applicability; one or more of the foregoing steps are performed in, by or for one or more processor devices.

41. (Currently Amended) A system for allocating a resource in a network comprising:

Art Unit: 2444

multiple network data ports to connect to the network to receive and route requests to servers in the network; first logic coupled to the network data ports for determining if any of a plurality of persistence policies, comprising at least one cookie-based persistence policy, at least one session-based persistence policy, and at least one client-based persistence policy, is applicable to a service request from a client in the network, and, if so, allocating the resource to the request based on application of the persistence policy determined to be applicable, wherein the first persistence policy, if applicable, directs the service request to a server, currently or previously connected to the client, that is identified based on content of the service request, the content comprising server, session or cookie information within the service request or corresponding packet at one or more layers corresponding to OSI layers 5-7, and wherein the second persistence policy, if applicable, directs the service request to a server, currently or previously connected to the client, that is identified based on client information within the service request or corresponding packet at a layer corresponding to OSI layer 4; the first logic configured so that: when a service request for a content-enabled service, comprising a request for a connection to a server at a layer corresponding to any of OSI layers 5-7, is received, the first logic determines if either the at least one cookie-based or at least one session-based persistence policy is applicable; when a service request for a non-content-enabled service, comprising a request for a connection to a server at a layer corresponding to OSI layer 4, is received, the first logic determines if the at least one client-based persistence policy is applicable; when the service request for a content-enabled service is received but it is determined that both the at least one client-based persistence policy and the at least one session-based persistence policy are inapplicable, the first logic determines if the at least

one client-based persistence policy is applicable; wherein the first logic is shared by and supports service requests for both content-enabled and non-content enabled services; and second logic in the network for allocating the resource to the request based on application of a load balancing policy only if none of the plurality of persistence policies is determined to be applicable as determined by the first logic wherein, when the service request for a content-enabled service is bound to a server identified through application of a load balancing policy, such request is bound to the server only after multiple persistence policies are checked for applicability; and one or more of the foregoing elements are implemented, embodied, maintained or controlled by one or more processor devices.

44.(New) The system of any one of claims 1, 3, and 41 wherein the one or more processor devices comprise one or more hardware processor devices, one or more processors executing software, or any combination of the foregoing.

45. (New) The system of claim 44 wherein the one or more hardware processor devices comprise one or more finite state machines.

46.(New) The system of claim 45 wherein the one or more finite state machines are implemented as one or more application specific integrated circuits (ASICs).

47. (New) THE method of any one of claims 20 and 21 wherein the one or more processor devices comprise one or more hardware processor devices, one or more processors executing software, or any combination of the foregoing.

48.(New) The method of claim 47 wherein the one or more hardware processor devices comprise one or more finite state machines.



49. (New) The method of claim 48 wherein the one or more finite state machines are implemented as one or more application specific integrated circuits (ASICs).

**Cancelled claims 5-19 and 23-38.**

***Reasons for Allowance***

The following is an examiner's statement of reasons for allowance: The prior art of record fails to teach wherein a service request for a content-enabled service, comprising a request for a connection to a server at a layer corresponding to any of OSI layers 5-7, is received, determining if the at least one first persistence policy is applicable; when a service request for a non-content-enabled service, comprising a request for a connection to a server at a layer corresponding to OSI layer 4, is received, determining if the at least one second persistence policy is applicable in combination with wherein the service request for a content-enabled service is bound to a server identified through application of a load balancing policy, such request is bound to the server only after multiple persistence policies are checked for applicability in combination with the limitation.

Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue fee. Such submissions should be clearly labeled "Comments on Statement of Reasons for Allowance."

Any inquiry concerning this communication or earlier communications from the examiner should be directed to DJENANE M. BAYARD whose telephone number is (571)272-3878. The examiner can normally be reached on Monday- Friday 5:30 AM- 3:00 PM.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, William Jr. Vaughn can be reached on (571) 272-3922. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/D. M. B./  
Examiner, Art Unit 2444  
/William C. Vaughn, Jr./  
Supervisory Patent Examiner, Art Unit 2444